

FRXLIB-0299467.01-JTENG
August 2, 2004 2:34 PM

LIST OF RELEVANT REFERENCES

The search revealed the following U.S. patents and patent applications, copies of which are listed for convenience:

<u>U.S. Patent No.</u>	<u>Inventor</u>
6,038,641	Zangenehpour
6,275,897 B1	Bachmat
6,470,419 B2	Take et al.

<u>U.S. Patent Application Publication No.</u>	<u>Inventor</u>
2003/0188104 A1	Sullivan
2004/0019740 A1	Nakayama et al.
2004/0059870 A1	Ash et al.
2004/0083338 A1	Moriwaki et al.

Discussion of References:

U.S. Pat. No. 6,275,897 B1 to **Bachmat** utilizes a remote slave mass storage subsystem as a cache for a mirrored master mass storage subsystem (Abstract). When a host computer wishes to retrieve information, if the information is not in the cache of the master mass storage subsystem, and if the information is still in the cache of the slave mass storage subsystem, the host computer will enable the information to be transferred from the slave mass storage subsystem to the master mass storage subsystem for access by the host computer (Fig. 1; col. 3, lines 11-16). However, the slave mass storage subsystem 12S is connected to the master mass storage subsystem 12M via a network communication link 16, rather than the internal bus 13 (~ internal connection section) connecting all together the communications interface section 15, the storage device interface section 21, 23, and the cache memory section 31. As such, **Bachmat** does not provide “a storage device control unit with an internal connection section for connecting all together a communications interface section, a storage device interface section, and a cache memory section with a first cache memory, while the connection section is provided with second cache memory for storage of data same as the data stored in the first cache memory” as now recited in claim 1 of the present invention.

U.S. Pat. Pub. No. 2004/0019740 A1 by **Nakayama** et al. assigned to Hitachi, Ltd. provides a destaging method for a storage apparatus system, i.e., a method of transferring to a first/lower hierarchy cache memory data stored in a second/higher hierarchy cache memory and scheduled to be written onto a storage unit, only “*when an interruption of a data storage function of the disk control apparatus for the information processing apparatus occurs (abstract)*”. “*When the client computer 100 shuts down, for example, its OS (operating system) begins a flush processing of dirty data stored in the corresponding cache 103 (S4000). The dirty data is transferred to the cache 203 of one of the NAS engines 200 via the LAN 500. When the flush processing is completed for all dirty data (S4001), the client computer 100 sends a flush completion notice to the NAS engine 200 and terminates the processing (S4002).*”[0055] As shown in Fig. 1, although the NAS server has two NAS engines each with a cache 203, the two caches 203 are indirectly connected with the LAN 500, rather than any internal connection section. Moreover, they do not cache identical data in each other. In addition, the first/lower hierarchy cache memory 203 and the second/higher hierarchy cache memory 103 are connected via LAN 500, rather than any internal connection section. **Nakayama** simply fails to teach any storage device with a control unit internal connection section (1) having one addition cache memory of it own and (2) connecting all together a communications interface section, a storage device interface section, and the cache memory 203 therein. As such, **Nakayama** does not provide “a storage device control unit with an internal connection section for connecting all together a communications interface section, a storage device interface section, and a cache memory section with a first cache memory, while the connection section is provided with second cache memory for storage of data same as the data stored in the first cache memory” as now recited in claim 1 of the present invention.

U.S. Pat. Pub. No. 2004/0059870 A1 by **Ash** et al. retrieves data backed-up in a non-volatile storage (NVS) 26b of a redundant cluster 20b to restore data cached in a cache 24a of a failed cluster 20a. (Fig. 1; [0004], [0024]). “*The NVS 26a, 26b in one cluster 20a, 20b is used to backup write data in the cache 24b, 24a in the other cluster 20b, 20a, e.g., NVS 26a backs up write data in cache 24b [0022].*” As such, none of the four components: the cache 24b, 24a and the NVSs 26a, 26b is dispensable for the **Ash** system. On the other hand, the invention only requires two caches for the minimum structure of the storage device control unit (broken lines in Fig. 1; p. 9, lines 10-12). Even if comparing only the

cache 24a, NVS 26b with the two caches of the invention, the NVS 26b is merely back-up memory accessed in the event of failure or loss of the original data, rather than a cache, i.e., a frequently accessed small fast memory holding recently accessed data and designed to speed up subsequent access to the same data. Further, even if comparing only the caches 24a, 24b with the two caches of the invention, the caches 24a, 24b simply do not cache identical data in each other. As such, **Ash** does not provide “a storage device control unit with an internal connection section for connecting all together a communications interface section, a storage device interface section, and a cache memory section with a first cache memory, while the connection section is provided with second cache memory for storage of data same as the data stored in the first cache memory” as now recited in claim 1 of the present invention.

U.S. Pat. Pub. No. 2004/0083338 A1 by **Moriwaki** et al. assigned to Hitachi, Ltd. provides for a disk array controller with two inter-connection networks. A first inter-connection network consisting of a plurality of connecting switches 2 to connect a plurality of disk array control units 1-1 is used for the connection between the channel IF units 11/disk IF units 12 and the shared memory units 13, and a second inter-connection network, consisting of a plurality of connecting switches 2 to connect a plurality of disk array control units 1-1, is used for the connection between the channel IF units 11/disk IF units 12 and the cache memory units 14 (Fig. 1; [0044]). **Moriwaki** does not teach or suggest any two of the cache memory units 14 cache identical data in each other. The system has a first interface with a host computer, a second interface unit having an interface with a first memory unit (i.e., the shared memory units 13), a second memory (i.e., the cache memory units 14) for storing data read/to be written from/in the first memory temporarily, and a transfer path for connecting the first/second interface unit to the second memory ([0020]). As such, **Moriwaki** does not provide “a storage device control unit with an internal connection section for connecting all together a communications interface section, a storage device interface section, and a cache memory section with a first cache memory, while the connection section is provided with second cache memory for storage of data same as the data stored in the first cache memory” as now recited in claim 1 of the present invention.

U.S. Pat. No. 6,038,641 to **Zangenehpour** divides a cache memory in a computer of the first embodiment (Fig. 1; Abstract) into two parts: a master cache memory and a

slave cache memory. The slave cache memory receives data from a slow source and then sends it quickly to the master cache memory so that the master cache memory is occupied less of the time with reading or writing data to or from a mass-storage or other peripheral device and is available more of the time for reading or writing data to or from main memory and responding to CPU requests. The first embodiment does not involve any other computer connected via a network. The second embodiment (Fig.3) applies the cache design into each of three work stations 30, 35, 36 connected via a network 34 to a central file server 33. There is simply no standing-alone storage device control unit (including a communications interface section, a storage device interface section, and a first cache memory section connected by an internal connection section with a second cache memory) placed between the central file server 33 and the work stations in Fig. 3. Rather, the central file server 33 has its own network interface and caches, and each of the work stations has its own network interface and caches. In other words, the master and slave caches are in a storage device 33 or a host computer 30, rather than in a standing-alone storage device control unit. Each of the work stations only handles an internal request for data stored externally, rather than handling data request for another host computer (col. 4 lines 51-53). As such, **Zangenehpour** does not provide “a storage device control unit with an internal connection section for connecting all together a communications interface section, a storage device interface section, and a cache memory section with a first cache memory, while the connection section is provided with second cache memory for storage of data same as the data stored in the first cache memory” as now recited in claim 1 of the present invention.

U.S. Pat. No. 6,470,419 B2 to **Take et al.** distributes and allocates the data for management of a page to separate caches (Abstract). The caches 16 (Fig. 2) simply do not cache identical data in each other. As such, **Take** does not provide “a storage device control unit with an internal connection section for connecting all together a communications interface section, a storage device interface section, and a cache memory section with a first cache memory, while the connection section is provided with second cache memory for storage of data same as the data stored in the first cache memory” as now recited in claim 1 of the present invention.

U.S. Pat. Pub. No. 2003/0188104 A1 by **Sullivan** provides a multi-level caching scheme (Fig. 1). A host cache 34 in the host system 12 is a level 1 cache, and the disk

cache 20 in the storage system 14 is a level 2 cache ([0020]). Although the data stored in the host cache 34 is also stored in the disk cache 20, the disk cache 20 also stores additional data not available in the host cache 34, rather than storing only data same as those in the host cache 34. Beside, there is simply no standing-alone storage device control unit (including a communications interface section, a storage device interface section, and a first cache memory section connected by an internal connection section with a second cache memory) placed between the host system 12 and the storage system 14 in Fig. 1. As such, **Sullivan** does not provide “a storage device control unit with an internal connection section for connecting all together a communications interface section, a storage device interface section, and a cache memory section with a first cache memory, while the connection section is provided with second cache memory for storage of data same as the data stored in the first cache memory” as now recited in claim 1 of the present invention.

CONCLUSION

Based on the results of the comprehensive prior art search as discussed above, Applicant contends that the position calculation method as now recited in independent claims 1 and 15, especially the features of “a double (or more levels) cache hierarchy construction” are patentably distinct from the cited prior art references.

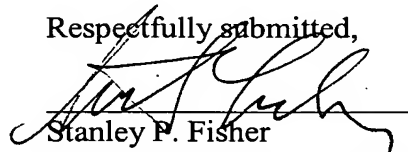
In particular, the storage device control unit 100 (e.g., Fig. 3) for performing data writing/reading to/from a storage device 300 responding to a request coming from an information processor 200 via a network of the invention, as recited in claim 1, comprises: a communications interface section 500 for carrying out communications with the information processor 200; a storage device interface section 600 for carrying out communications with the storage device 300; a cache memory section 700 including first cache memory 710 for storage of data coming and going between the information processor 200 and the storage device 300; and an internal connection section 800 for connecting all together the communications interface section 500, the storage device interface section 600, and the cache memory section 700 for communications thereamong. The connection section 800 is provided with second cache memory 820 for storage of data same as the data stored in the first cache memory 700. Such an internal double (or more levels) cache hierarchy construction (p. 14, last line) allows flexible data arrangement depending on the data contents to increase the cache hit ratio of the storage device control unit 100 in its entirety (p. 15, lines 6-9).

Claim 15 recites a method for controlling communications between an information processor and a storage device via a network comprising: providing a storage device control unit including a communications interface section for carrying out communications with an information processor, a storage device interface section for carrying out communications with a storage device, a cache memory section including first cache memory, and an internal connection section for connecting all together the communications interface section, the storage device interface section, and the cache memory section for communications thereamong, and the connection section being provided with second cache memory; storing data coming and going between the information processor and the storage device in the first cache memory; and storing in the second cache memory data same as the data stored in the first cache memory.

In view of all the above, clear and distinct differences as discussed exist between the present invention as now claimed and the prior art references, Applicant respectfully contends that the prior art references cannot anticipate the present invention or render the present invention obvious. Rather, the present invention as a whole is distinguishable, and thereby allowable over the prior art.

Favorable consideration of this application is respectfully solicited. Should there be any outstanding issues requiring discussion that would further the prosecution and allowance of the above-captioned application, the Examiner is invited to contact the Applicant's undersigned representative at the address and phone number indicated below.

Respectfully submitted,



Stanley P. Fisher
Registration Number 24,344

Juan Carlos A. Marquez
Registration Number 34,072

REED SMITH LLP
3110 Fairview Park Drive
Suite 1400
Falls Church, Virginia 22042
(703) 641-4200

August 2, 2004

SPF/JCM/JT